

Department of Aerospace Engineering
and Engineering Mechanics
University of Cincinnati

Groundspeed filtering for CTAS

11/26/94
10-35-22
0 217
30651
168

G.L. Slater
Department of Aerospace Engineering
and Engineering Mechanics
University of Cincinnati
Cincinnati, OH 45221

N95-15785
Unclas

G3/35 0030651

Nov. 11, 1994

Work done under NASA Cooperative Agreement NCC2-669
NASA Ames Research Center

(NASA-CR-197223) GROUNDSPD
FILTERING FOR CTAS (Cincinnati
Univ.) 16 p

1 Introduction

This report is the first part of our report for the 1993-94 year work on NASA NCC2-669, a cooperative agreement between NASA Ames, Air Traffic Management Branch and the University of Cincinnati, Gary L. Slater, Principal Investigator. This current report analyzes the current algorithm for filtering radar ground speed and proposes an alternate filtering algorithm from that currently used in CTAS. A second report authored by University of Cincinnati graduate student Michael Bolender is bound separately and looks at observed characteristics of CTAS trajectories and the predictive ability of CTAS. A third effort to look at scheduling and runway allocation for arrivals at airports with multiple runways is being continued and we expect our report on that to be available in about one month.

2 Groundspeed and speed filtering

2.1 Current status

Ground speed is one of the radar observables which is obtained along with position and heading from Center radar. Within CTAS, groundspeed is converted into airspeed using the wind speeds which CTAS obtains from the NOAA weather grid. This airspeed is then used in the trajectory synthesis logic which computes the trajectory for each individual aircraft. The time history of the typical radar groundspeed data is generally quite noisy, with high frequency variations on the order of five knots, and occasional “outlier’s” which can be significantly different from the probable true speed. To try to smooth out these speeds and make the ETA estimate less erratic, filtering of the ground speed is done within CTAS. In its base form, the CTAS filter is a “moving average” filter which averages the last ten radar values. In addition, there is separate logic to detect and correct for “outliers”, and acceleration logic which limits the groundspeed change in adjacent time samples. As will be shown, these additional modifications do cause significant changes in the actual groundspeed filter output.

2.2 Groundspeed filter results

To consider some of the actual observed behavior of groundspeed we looked at several aircraft in several data sets. The primary aircraft shown here is a turbojet from the May 11, 1994 dataset at the Dallas Fort Worth airport. This is flight AAL154, a Boeing 757 type of aircraft. A plot of the raw groundspeed and CTAS filtered ground speeds are shown in Figure 1. This plot is fairly typical in that several features are present:

1. There is an initial transient where the filtered airspeed dips away from the actual groundspeed before tracking the observed data.

2. The filtered groundspeed lags behind the observed groundspeed during any period of changing groundspeed. This may be in a constant altitude acceleration, but more typically is seen in a constant CAS or Mach descent.
3. During periods of constant groundspeed, the filter acts to smooth out the observed variations very well.

One feature immediately noted is that the current filter lags behind the airspeed a significant amount. For AAL154 the lag is 30-40 kts over a substantial time during the descent. This seems to be typical of a number of aircraft we have observed. Since a pure moving average filter will average the last N points, the average will lag behind in a constant acceleration by $N/2$ points. (Actually the outlier logic in CTAS always throws away one extra point, which in this case can cause the lag to be $(N + 1)/2$ points.) In CTAS the number of averages is set by `MAX_GS_AVERAGING` (in `cm_defs.h`) which is currently set to 10. For the observed deceleration this lag should be only about 25-30 knots (decelerations of five knots per data point are common.). While this magnitude of lag is observed, there were discrepancies in a comparison with a true 10 point moving average filter, indicating that the acceleration and/or outlier logic is having a significant effect on the ground speed estimate.

The current outlier logic in CTAS can in fact worsen the acceleration error through its outlier logic. First, if a point is outside of the acceleration limits set by CTAS (nominally $0.1g$), the actual data point is thrown away and replaced by a value at the acceleration limit. Secondly, the outlier logic always throws away one point of the last ten groundspeed values. (The point thrown away is the largest magnitude deviation from the ten point average.) When the aircraft is accelerating, as in a descent, the worst point is often the current point, meaning that the current groundspeed measurement is ignored and the previous nine points are averaged.

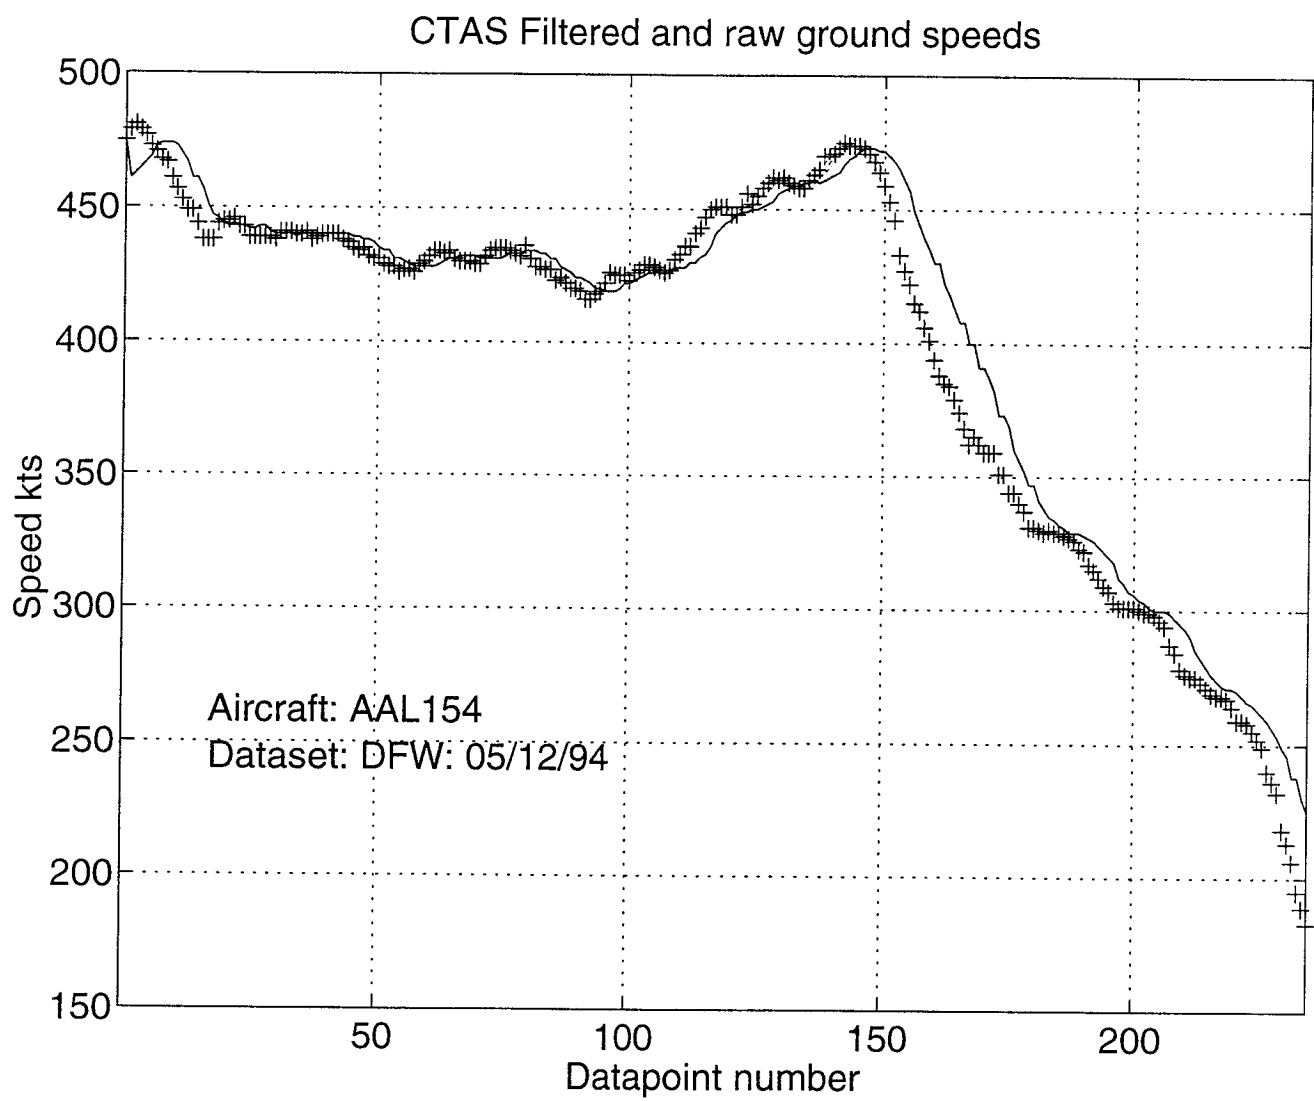


Figure 1: Raw and filtered groundspeed for AAL154

A final point is that the initial transient observed in the groundspeed estimate is caused by a logic error in the filtering code. In the initialization of the filter, the array that holds the last ten averages is filled with the flight plan true airspeed. (Nominally the planned true airspeed at cruise.) The first filtered ground speed is taken as the actual first data point, but all subsequent filtered output values average ten points using the filed speed as data. Because of the outlier and acceleration logic, the error caused by this propagates well beyond the first ten data points. Because the filed true airspeed is generally significant different from the groundspeed, this produces the large spike observed in the measured groundspeed, and the corresponding spike in the ETA estimate.

To help see exactly what is happening, and what can be done to improve filter estimates, several filter variants were considered. First an eleven point “smoother” was used to determine a smoothed estimate (using five points ahead of and behind to compute a centered average. This is used as a reference to test for the noise in the data. Figure 2 shows: (1) The error in the unfiltered ground speed, (2) the error in the CTAS filter, and (3) the error in a simple moving average filter using the last ten points. The groundspeed rate observed should not trigger the maximum rate constraint of 1.90626 kts/sec, but it is clear that the CTAS filter has additional errors over an averaging filter with no outlier removal. Note that the startup logic causes a large error in the initial groundspeed.

2.2.1 Kalman filter averaging filter

Because the pure moving average has an inherent delay, it is unable to track closely the groundspeed during any period of speed change. To improve the tracking capability, we considered an alternate filter which assumes the aircraft can have an acceleration . Even though the data points do not have a strict time tag associated with each point (the exact time of the radar hit is not known) a discrete Kalman filter was used to model the transition dynamics. The transition model uses acceleration and velocity as states, with acceleration modeled as a constant driven by white noise. Ground speed is used as the measurement,

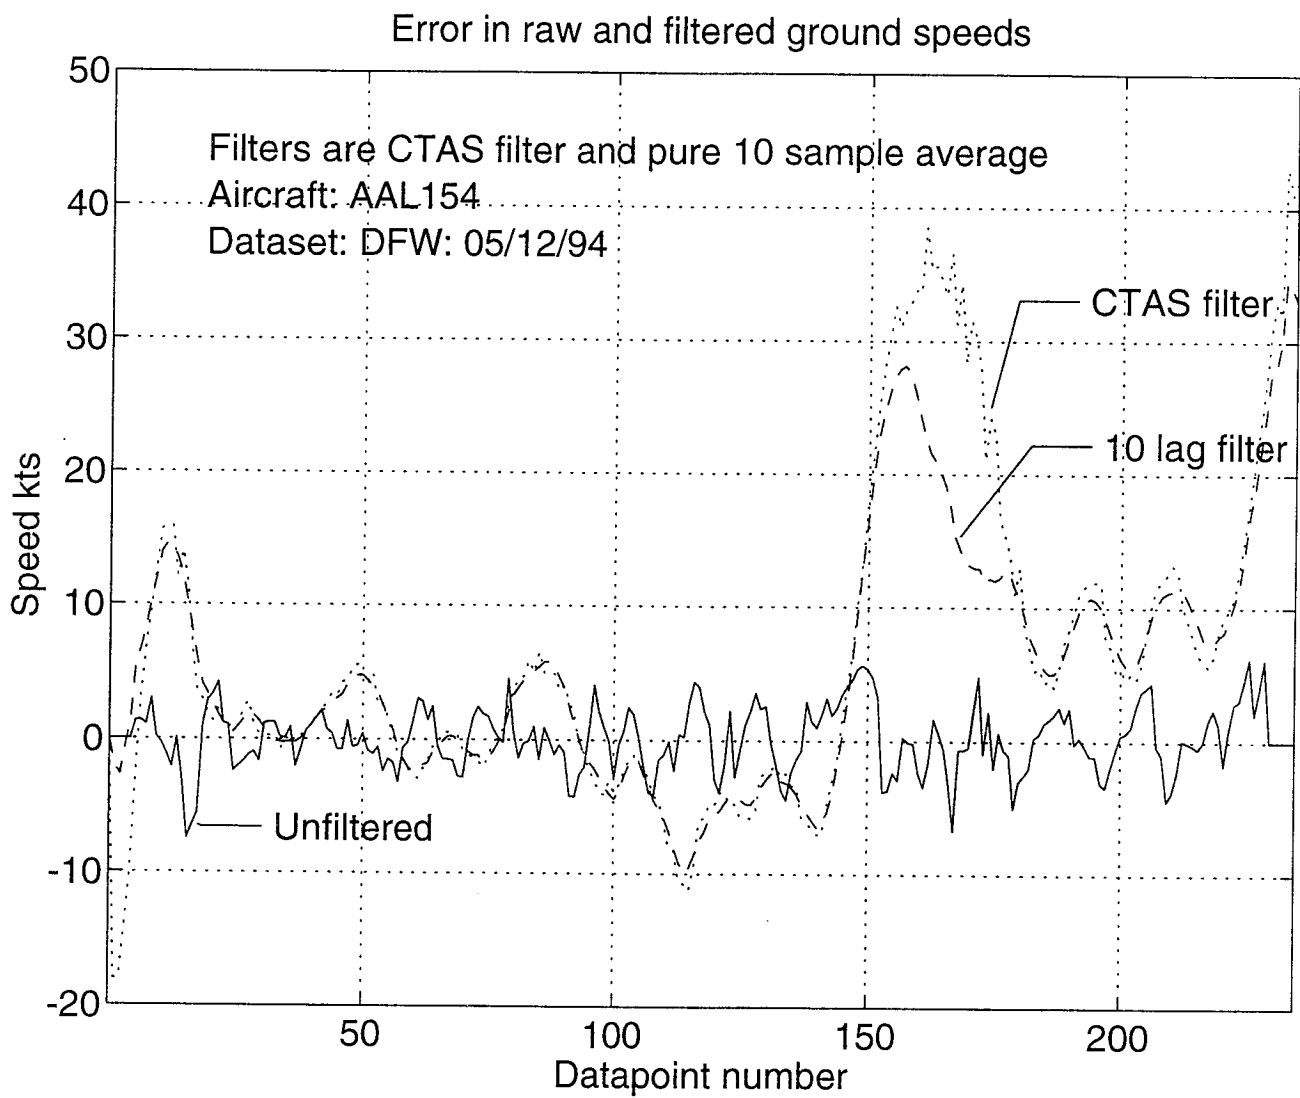


Figure 2: Raw and filtered groundspeed for AAL154

and the process and measurement noise covariances were adjusted to give the *appearance* of good filtering (determining by visually inspecting the data). We tested both a time varying filter and a steady state filter. There was a slight difference in the first few value of ground speed, but virtually identical after this starting transient. For this reason we implement the steady state filter saving significantly on the computations required to implement this filter.

From a mathematical point of view, the state vector is defined at time i to be $x_i = [a_i \ v_i]^T$, with the transition relation

$$\begin{aligned} x_{i+1} &= \Phi x_i + \Gamma w_i \\ y_i &= H x_i + r_i \end{aligned} \tag{1}$$

where y_i is the unfiltered ground speed and

$$\Phi = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \quad \Gamma = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad H = \begin{bmatrix} 0 & 1 \end{bmatrix}$$

The Kalman filter is implemented as

$$\begin{aligned} \hat{x}_i^{(-)} &= \Phi \hat{x}_{i-1}^{(+)} \\ \hat{x}_i^{(+)} &= \hat{x}_i^{(-)} + K_F (y_i - H \hat{x}_i^{(-)}) \\ \hat{y}_i &= H \hat{x}_i^{(+)} \end{aligned} \tag{2}$$

where y_i is the unfiltered ground speed, \hat{y}_i is the Kalman filtered ground speed, and K_F is the Kalman gain matrix.

The noise covariance for the process and measurement noise were adjusted strictly to give “reasonable” transient response of the filter. For the filter results presented here we used $W = E(w_i^2) = 0.01$ and $V = E(r_i^2) = 1$. The closed loop filter eigenvalues are complex at $\zeta = 0.82 \pm j0.22$, giving a “settling time” of about 15 samples. The filter appears to give good results and tracks the velocity much closer on acceleration/deceleration segments than the current filter, while having about the same performance on constant speed segments.

We did not add any extra “outlier” logic although that may be desirable. The data we examined had few or no points we would consider as true outliers. For the Kalman filter initialization, the first data point is used as measured; subsequent points are filtered using the filter logic.

We have not tested the Kalman filter code in CTAS but this would be a very simple code to implement. The current logic stores the previous 10 data points and does roughly 30 additions, two divides, and a number of logical comparisons. The Kalman logic only stores the previous state (dimension =2), the two constant gains, and uses only two multiplies and three adds. Additional logic would need to be added for outlier detection and prevention. The Kalman gain we used was the steady state filter gain which for the numbers given has the value $K = [0.0799 \ 0.3618]^T$.

The result of the Kalman filter applied to the same data previously cited is shown in Figures 3, 4.

The data shown for AAL154 is representative of several turbojet aircraft which we examined. Because turboprop aircraft exhibit much greater speed variations than the turbojets, we also tested the filter code on several aircraft of this type. Shown here is data for aircraft BTA2277 from a February 2, 1994 dataset at Denver Stapleton airport. (The aircraft type is listed as a BE02.) Shown below are the corresponding figures for this turboprop aircraft. As is expected the turboprop groundspeed exhibits considerably more variation than does the turbojet. The same Kalman filter tracks the speed variation quite well and seems to be an improvement over the current logic. Since these two classes of aircraft seem to have significantly different speed characteristics, there may be some merit to having different filter matrices for different classes of aircraft, giving different filter dynamics for turbojets and turboprops for example. Using the Kalman filter logic it would be very easy to program different filter dynamics for different classes of aircraft.

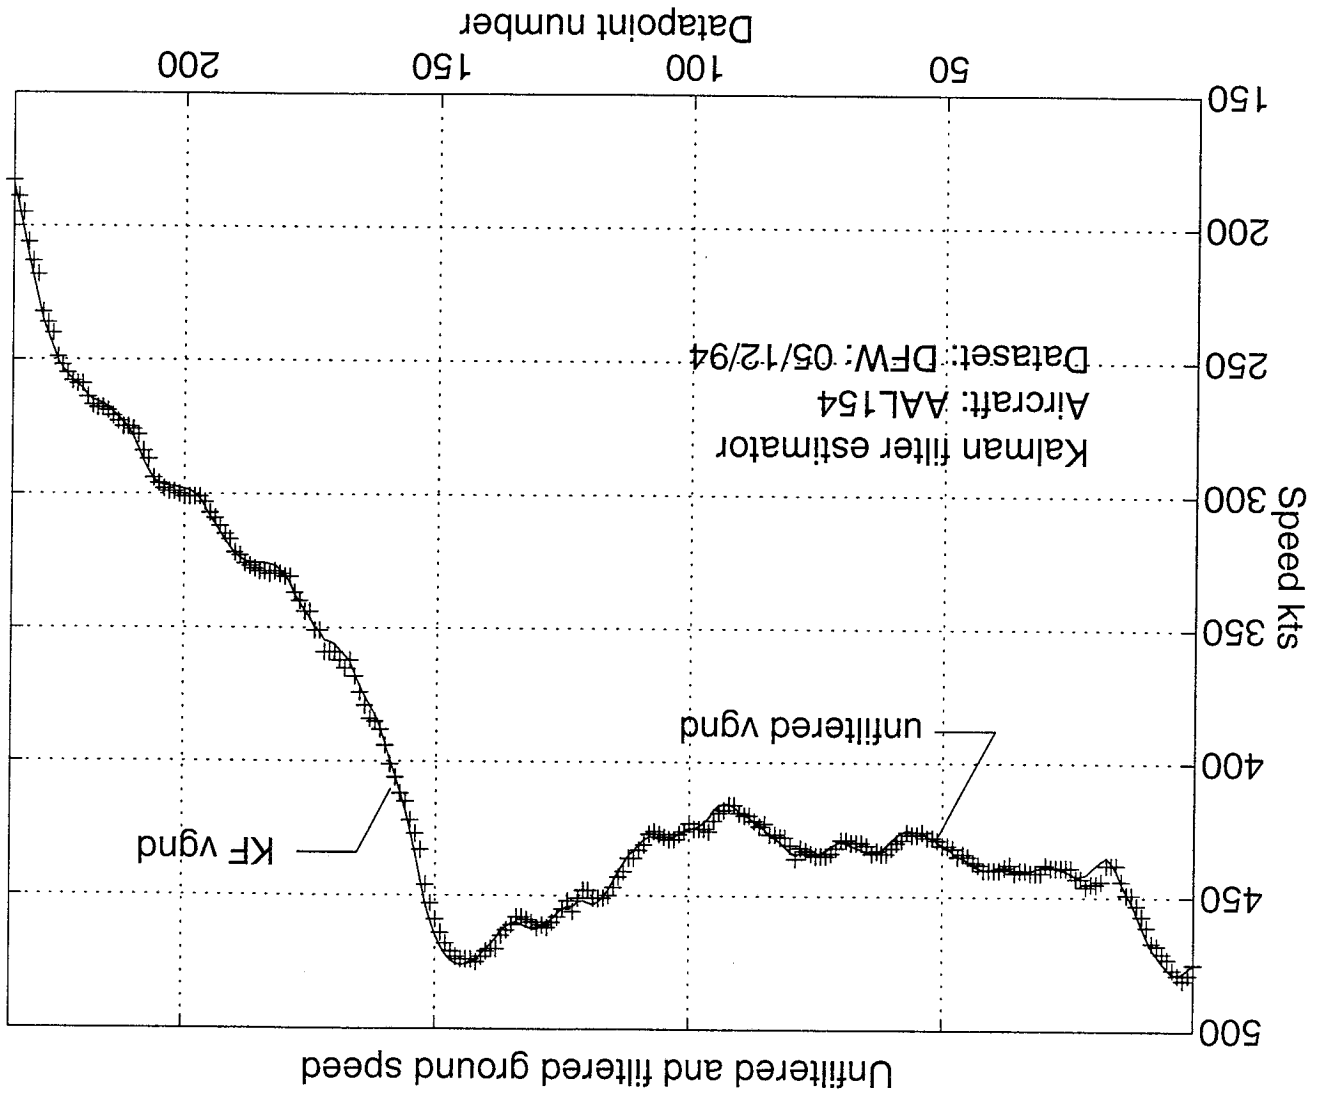


Figure 3: Kalman estimate and unfiltered groundspeed

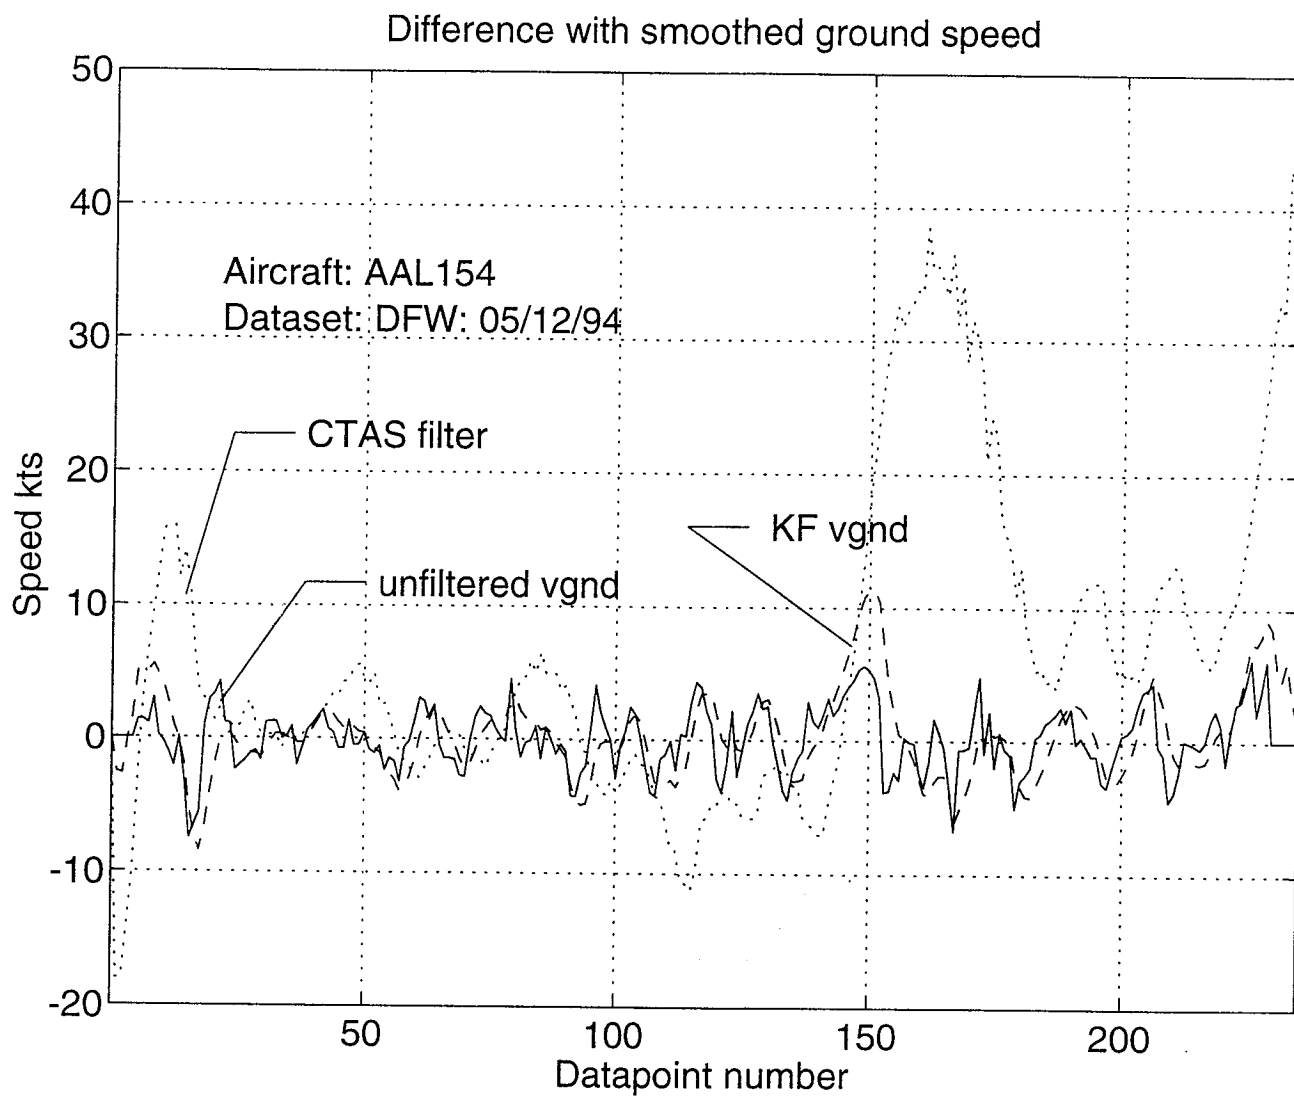


Figure 4: Difference in CTAS and Kalman filters on groundspeed

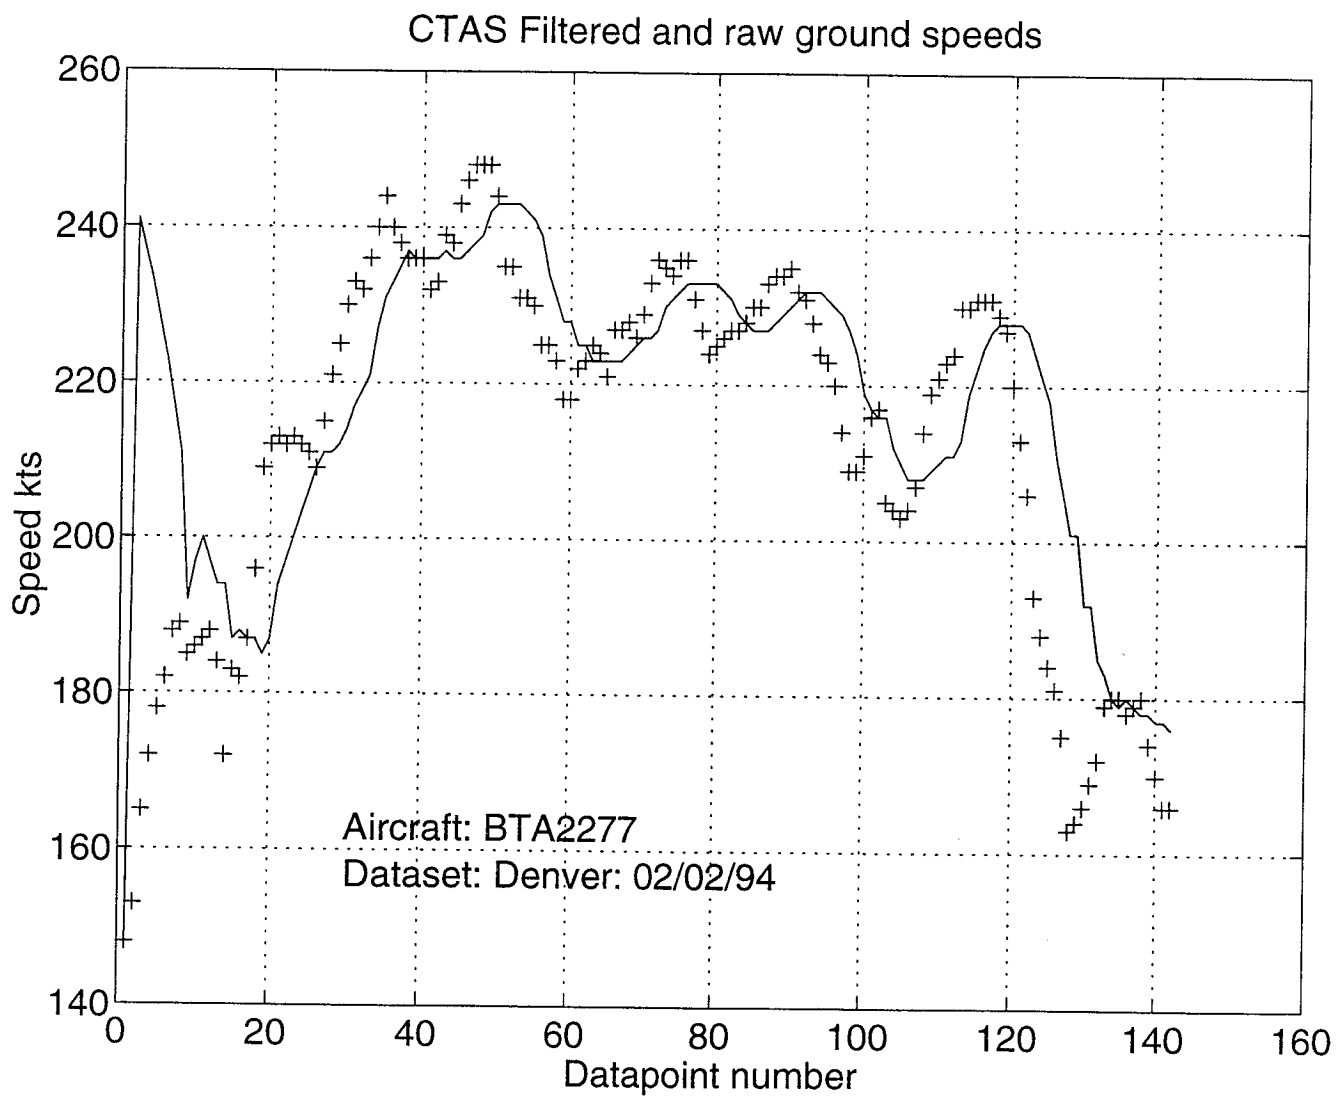


Figure 5: Raw and filtered groundspeed

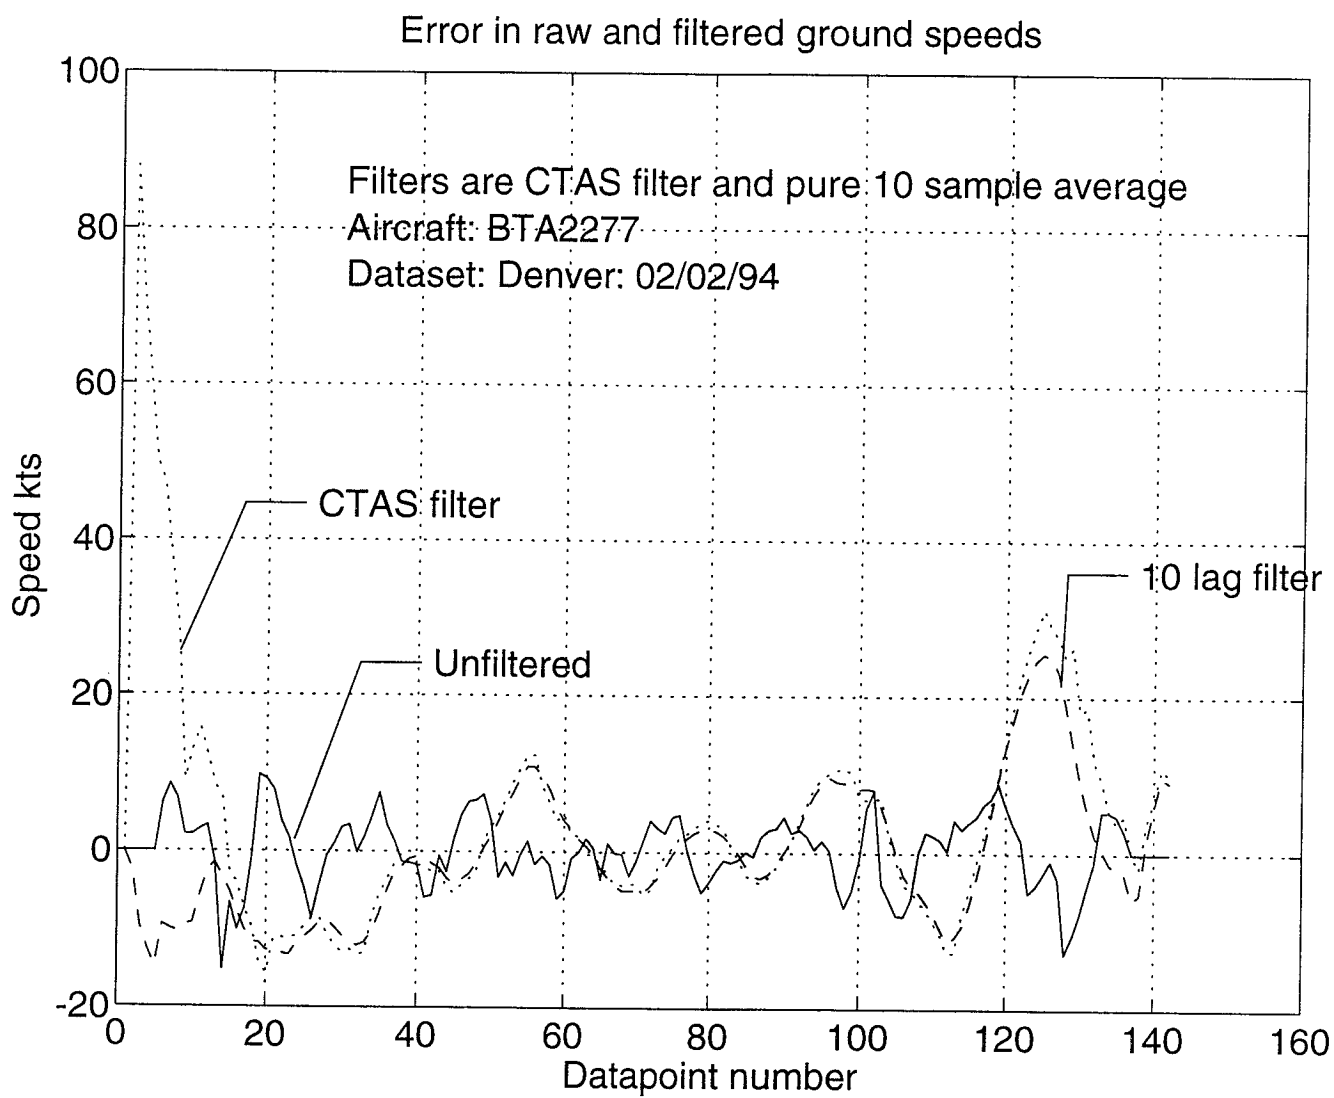


Figure 6: Groundspeed difference

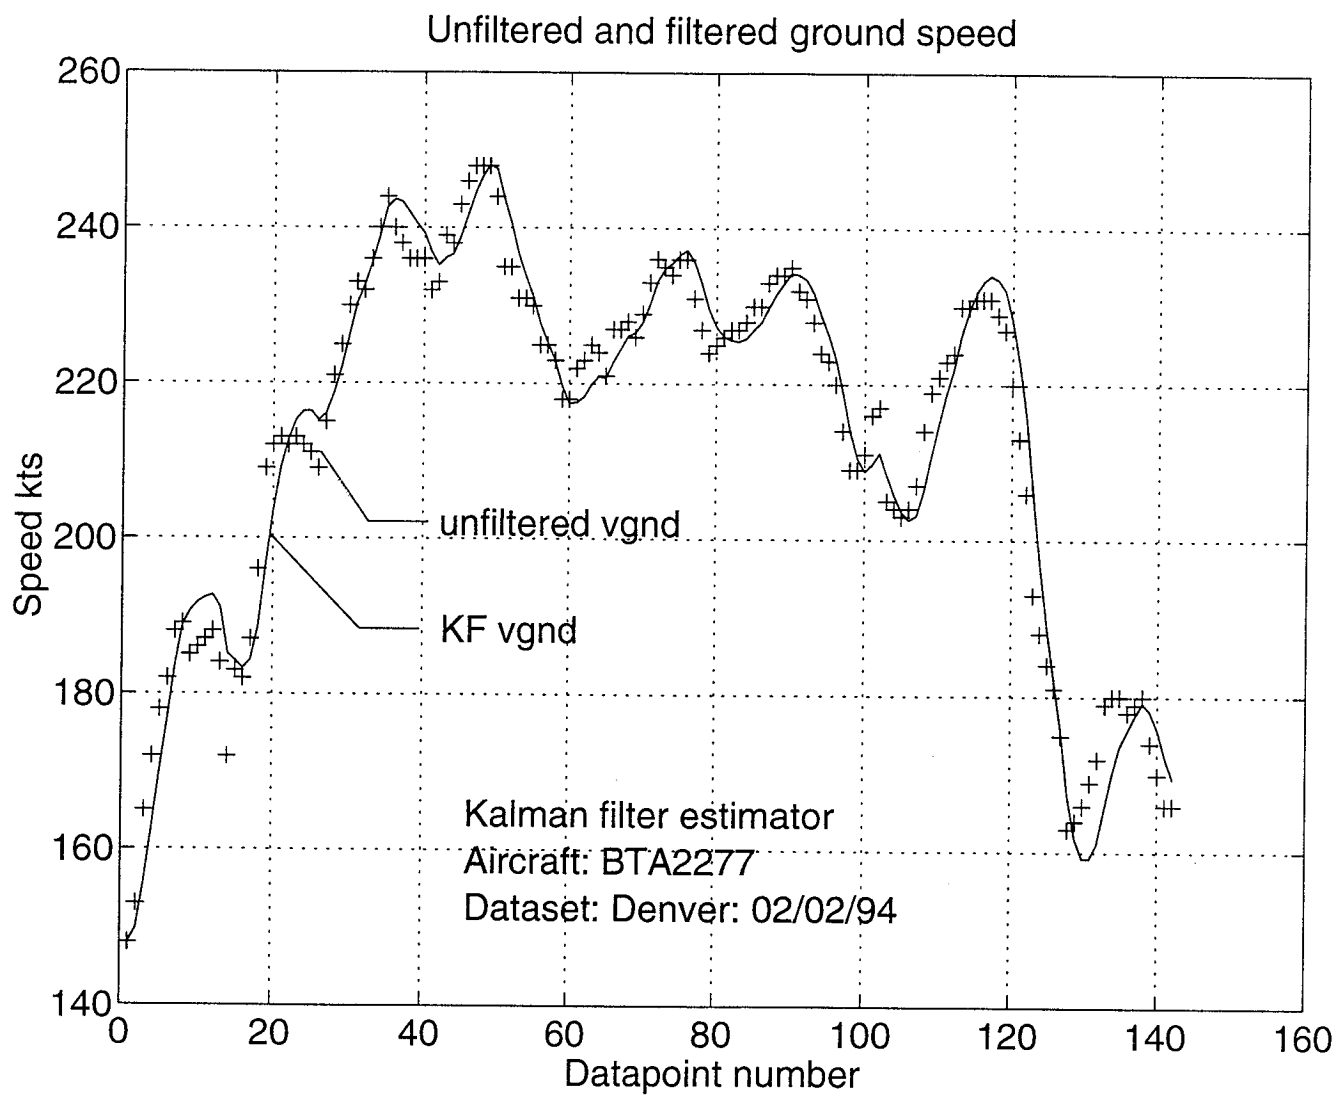


Figure 7: Kalman estimate and unfiltered groundspeed

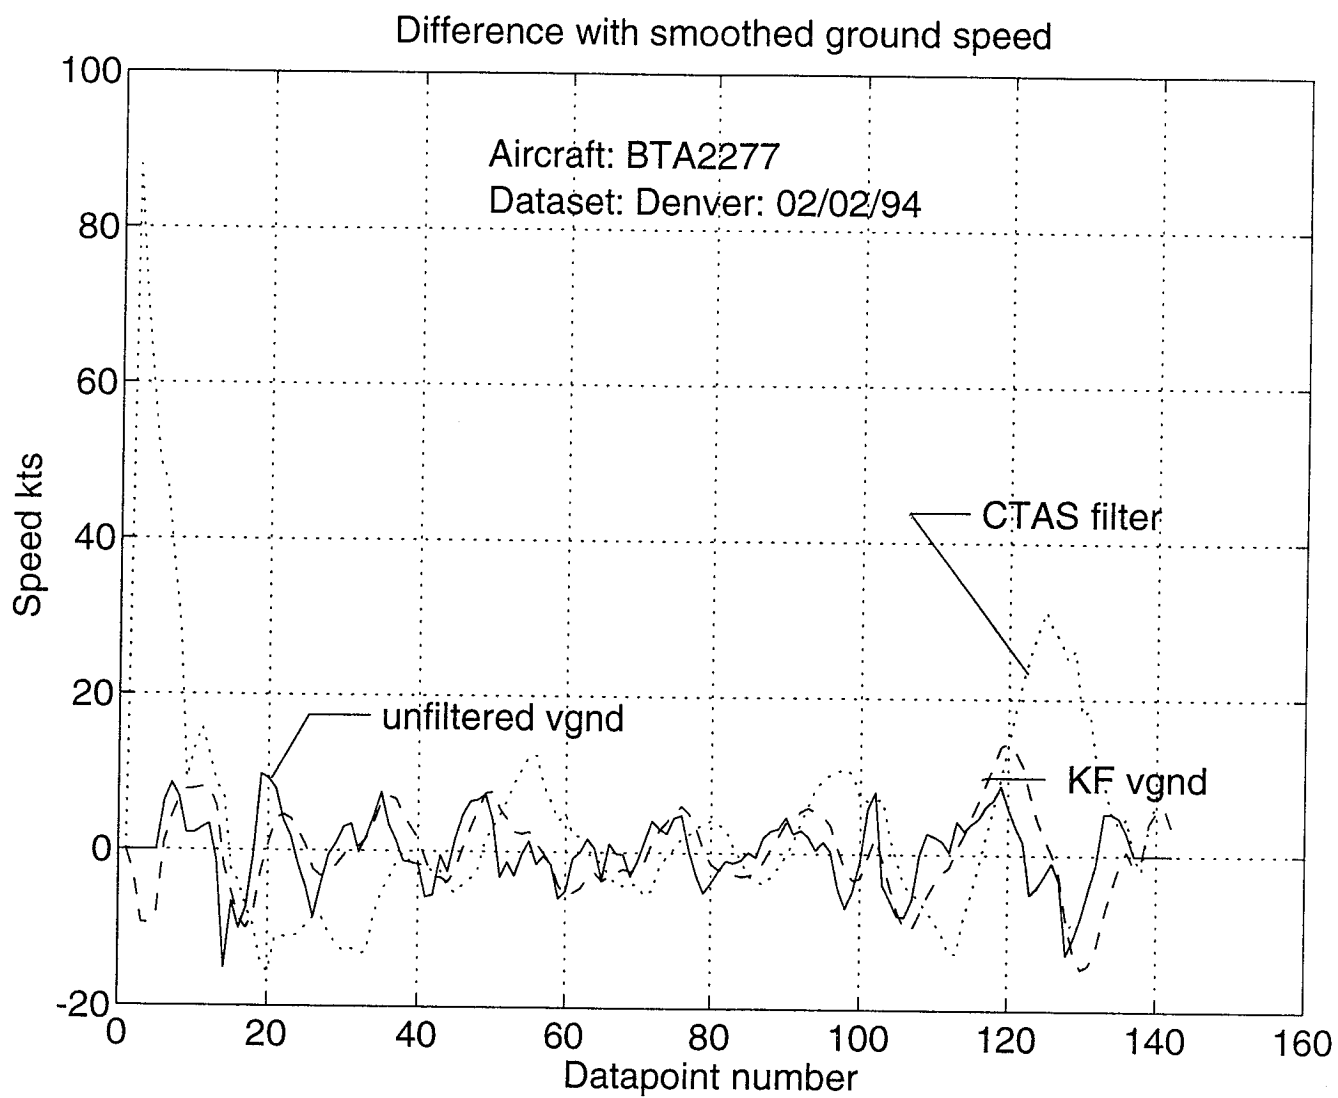


Figure 8: difference in filtered and smoothed groundspeed

2.3 Conclusion

Our conclusion is that the current ground speed filter logic is unable to track accurately the speed variations observed on many aircraft. The Kalman filter logic however, appears to be an improvement to the current algorithm used to smooth ground speed variations, while being simpler and more efficient to implement. Additional logic which can test for true "outliers" can easily be added by looking at the difference in the a priori and post priori Kalman estimates, and not updating if the difference in these quantities is too large. While not presented in this report (see accompanying volume), one test we use for consistency of CTAS data is to plot true airspeed versus altitude over a trajectory. The Kalman filter implementation seems to give a truer trajectory in that constant CAS and constant mach segments are more clearly visible during a descent.

While the velocity accuracy appears to be improved in these off-line tests, it is not clear at this time how this improvement will affect the accuracy of the ETA predictions, which are of course the primary goal of this effort. A test of this algorithm should be done within CTAS to quantify improvement in the time estimates.

We believe the ETA estimates are also affected to a lesser degree by the heading which exhibits the same (or worse) noise characteristics as the radar groundspeed. Currently the heading is not filtered at all because the heading changes quite frequently, and the moving average filter would be unable to track these changes accurately. We believe the proposed Kalman filter logic could also be effective in smoothing the heading data. We are currently undertaking a project to examine an improved implementation of the filtering algorithm presented here for groundspeed, and the possible extension of this logic to include filtering of position and heading using a kinematically consistent model. This extension will include an improved outlier detection and correction algorithm based on fuzzy logic to help distinguish between bad data and aircraft maneuvering. We hope to present the results of this investigation in the near future.